arXiv:1511.03225v3 [cs.LG] 29 Jul 2016

On the Geometry of Output-code Multi-class Learning

Maria Florina Balcan
School of Computer Science
Carnegie Mellon University
ninamf@cs.cmu.edu

Travis Dick
School of Computer Science
Carnegie Mellon University
tdick@cs.cmu.edu

Yishay Mansour
Microsoft Research and
Tel Aviv University
mansour@tau.ac.il

August 1, 2016

Abstract

We provide a new perspective on the popular multi-class algorithmic techniques one-vs-all and (error correcting) output-codes. We show that in cases where these techniques are successful (at learning from labeled data), they implicitly assume structure on how the classes are related. By making that structure explicit, we design algorithms to recover the classes based on limited labeled data. We provide results for the commonly studied cases of one-vs-all learning and when the code words of the classes are well separated. We additionally consider the more challenging case where the codewords are not well separated, but satisfy a boundary features condition that formalizes the natural assumption that every semantic feature should be significant.

1 Introduction

Motivation: Large scale multi-class learning problems with an abundance of unlabeled data are ubiquitous in modern machine learning. For example, an in-home assistive robot needs to learn to recognize common household objects, familiar faces, facial expressions, gestures, and so on in order to be useful. Such a robot can acquire large amounts of unlabeled training data simply by observing its surroundings, but it would be prohibitively time consuming (and frustrating) to ask its owner to annotate any significant portion of this raw data. More generally, in many modern learning problems we often have easy and cheap access to large quantities of unlabeled training data (e.g., on the internet) but obtaining high-quality labeled examples is relatively expensive. More examples include text understanding, recommendation systems, or wearable computing [Thrun, 1996, Thrun and Mitchell, 1995b,a, Mitchell et al., 2015]. The scarcity of labeled data is especially pronounced in problems with many classes, since supervised learning algorithms typically require labeled examples from every class. In such settings, algorithms should strive to make the best use of unlabeled data in order to minimize the need for expensive labeled examples.

Overview: We approach label-efficient learning by making the implicit assumptions of popular multi-class learning algorithms explicit and showing that they can be more fully exploited when learning from limited labeled data. We focus on a family of techniques called *output codes* that work by decomposing a given multi-class problem into a collection of binary classification tasks [Mohri et al., 2012, Dietterich and Bakiri, 1995, Langford and Beygelzimer, 2005, Beygelzimer et al., 2009]. The most well known output code technique is one-vs-all learning, where we learn one binary classifier for distinguishing each class from the union of the rest. When output codes are successful at learning from labeled data, it often implies geometric structure in the underlying problem. For example, if it is possible to learn an accurate one-vs-all classifier with linear separators, it implies that no three classes can be collinear, since then it would be impossible for a single linear separator to distinguish the middle class from the union of the others. In this work exploit this implicitly assumed structure to design label-efficient algorithms for the commonly assumed cases of one-vs-all and error correcting output codes, as well as a novel boundary features condition that captures the intuition that every semantic feature should be significant.

Our results: Before discussing our results, we briefly review the output code methodology. For a problem with L classes, a domain expert designs a code matrix $C \in \{\pm 1\}^{L \times m}$ where each column partitions the classes into two meaningful groups. For example, when recognizing household objects we could use the following true/false questions

to define the partitions: "is it made of wood?", "is it sharp?", "does it have legs?", "should I sit on it?", and so on. Each row of the code matrix describes one of the classes in terms of these partitions (or semantic features). For example, the class "table" could be described by the vector $(+1, -1, +1, -1)$, which is called the class's code word. Once the code matrix has been designed, we train an output code by learning a binary classifier for each of the binary partitions (e.g., predicting whether an object is made of wood or not). To predict the class of a new example, we predict its code word in $\{\pm 1\}^m$ and output the class with the nearest code word under the Hamming distance. Two popular special cases of output codes are one-vs-all learning, where C is the identity matrix (with -1 in the off-diagonal entries), and error correcting output codes, where the Hamming distance between the rows of C is large.

In each of our results we assume that there exists a consistent linear output code classifier and we impose constraints on the code matrix and the distribution that generates the data. We present algorithms and analysis techniques for a wide range of different conditions on the code matrix and data distribution to showcase the variety of implicit structures that can be exploited. For the code matrix, we consider the case when the code words are well separated (i.e., the output code is error correcting), the case of one-vs-all (where the code matrix is the identity), and a natural boundary features condition that intuitively requires every semantic feature to be significant. These conditions can loosely be compared in terms of the Hamming distance between code words. In the case of error correcting output codes, the distance between code words is large (at least $d + 1$ in a d-dimensional problem), in one-vs-all the distance is always exactly 2, and finally in the boundary features condition the distance can be as small as 1. In the latter cases, the lower distance is compensated by other structure in the code matrix. For the distribution, we either assume that the distribution is nearly uniform (i.e., it has a density function that is both upper and lower bounded on its support), or we assume that the distribution has a thick level set property, which essentially rules out distributions whose level sets have bridges or cusps that are too thin. Both regularity assumptions are used to establish a connection between the geometry implied by the consistent output code classifier and the data distribution.

One-vs-all: We first showcase how to exploit the implicit structure assumed by the classic linear one-vs-all approach when the data lies in the unit ball and the data distribution is nearly uniform. We show that in this case it is rare to see nearby points that belong to different classes, since the sets of points belonging to classes i and j can only touch at a point (where their corresponding linear separators cross). The key observation is that after projecting the data onto the surface of the unit ball, the existence of the consistent one-vs-all classifier allows us to give upper and lower bounds on the the density of the projected data. Using these bounds, we show that the density must go to zero at the boundaries of classes and that it is bounded below on the interior of the classes. Based on these observations, we design a label-efficient algorithm that learns by estimating the connected components of the ϵ level set of the projected density.

Error correcting output codes: We also consider the commonly used and natural case of linear output codes where the Hamming distance between code words is large. In practice, output codes are designed to have this property in order to be robust to prediction errors for the binary classification tasks. We suppose that the output code makes at most β bits of error during prediction and has code words with Hamming distance at least $2\beta+d+1$ in a d-dimensional problem. The key observation is that when the code words are well separated, this implies that points belonging to different classes must be geometrically separated as well. When the data distribution is nearly uniform, the points belonging to each class will also be well connected, so the classes can be learned by a simple threshold clustering algorithm. However, when the data distribution is not uniform, we are not guaranteed that all the points belonging to class i will form a single cluster. Under the much weaker condition that the data distribution satisfies a thick level set property (requiring that its level sets do not have bridges or cusps that are too thin), we use the existence of the consistent linear output code to show that there will be a pruning of the single linkage hierarchical clustering whose clusters are label homogeneous. In this case, our active learning algorithms query a small number of labels to find this good pruning.

The boundary features condition: Finally, we introduce a natural and interesting condition on the code matrix that captures the intuition that every semantic feature should be significant. This condition has the weakest separation requirement (i.e., code words may have Hamming distance 1). The boundary features condition requires that every semantic feature constrains the support of the data distribution. The key observation is that we can directly approximate the linear separators defining the consistent output code by looking for hyperplanes that locally separate regions of space with data from regions without. Unlike the previous two conditions, our algorithm for this setting is not based on clustering and, rather than approximating the set of points that belong to each class, it approximates the global

structure of the problem by finding the underlying linear separators.

Our results show that the existence of a consistent output code classifier can imply many different types of geometric constraints on the underlying learning problem, and that under mild distributional assumptions these geometric constraints can be exploited to learn from mostly unlabeled data.

2 Related Work

Reduction to binary classification is one of the most widely used techniques in applied machine learning for attacking multi-class problems. Indeed, the one-vs-all, one-vs-one, and the error correcting output code approaches [Dietterich and Bakiri, 1995] all follow this structure [Mohri et al., 2012, Langford and Beygelzimer, 2005, Beygelzimer et al., 2009].

There has been growing interest in learning from mostly unlabeled data in multi-class problems. Donmez et al. [2010] show that under certain modeling assumptions it is possible to derive consistent estimators for the risk of classifiers based purely on unlabeled data. Their modeling assumptions require that the marginal distribution over class labels is known and that the data is generated according to a known family of models (e.g., univariate linear regression, where we assume that $y = ax + \epsilon$ for $\epsilon \sim N(0, \tau^2)$). Assuming the data is generated from a known class of models is similar to our assumption that there exists a consistent output code classifier. Continuing this work, Balasubramanian et al. [2011a,b] show that for linear classification problems where the feature vectors x are high dimensional with weakly correlated features it is possible to consistently estimate several margin-based risks of a linear classifier from unlabeled data by exploiting the fact that the classifier score is distributed as a mixture of Gaussians. The motivation for this line of work is similar to our own, but their techniques are significantly different and we consider commonly used discriminative assumptions.

Another work related to ours is that of Balcan et al. [2013], where labels are recovered from unlabeled data. The main tool that they use, in order to recover the labels, is the assumption that there are multiple views and an underlying ontology that are known, and restrict the possible labeling. Similarly, Steinhardt and Liang [2016] show how to use the method of moments to estimate the risk of a model from unlabeled data under the assumption that the data has three independent views. Our work is more widely applicable, since it applies when we have only a single view.

The output-code formalism is also used in [Palatucci et al., 2009] for the purpose of zero shot learning. They demonstrate that it is possible to exploit the semantic relationships encoded in the code matrix to learn a classifier from labeled data that can predict accurately even classes that *did not appear in the training set*. These techniques make very similar assumptions to our work but require that the code matrix C is known and the problem that they solve is different.

3 Preliminaries

We consider multiclass learning problems over an instance space $\mathcal{X} \subset \mathbb{R}^d$ where each point is labeled by $f^* : \mathcal{X} \to \{1, \ldots, L\}$ to one out of L classes and the probability of observing each outcome $x \in \mathcal{X}$ is determined by a data distribution P on $\mathcal{X}$. In all of our results we assume that there exists a consistent (but unknown) linear output-code classifier defined by a code matrix $C \in \{\pm 1\}^{L \times m}$ and m linear separators $h_i(x) = w_i^\top x - b_i$ for $i = 1, \ldots, m$. Without loss of generality we assume that $\|w_i\| = 1$. We denote class i's code word by C_i and define $h(x) = (\mathrm{sign}(h_1(x)), \ldots, \mathrm{sign}(h_m(x)))$ to be the predicted code word for point x.

Our goal is to learn a hypothesis $\hat{f} : \mathcal{X} \to \{1, \ldots, L\}$ from an unlabeled sample drawn from the data distribution P, possibly using a small amount of labeled data. Some of our algorithms learn from only unlabeled data. In these cases, since it is impossible to recover the label identities from an unlabeled sample, our goal is to minimize the following permutation-invariant error criterion: $\mathrm{err}_P(\hat{f}) = \min_{\sigma \in \mathbb{S}} \mathbb{P}_{x \sim P}\big(\sigma(\hat{f}(x)) \neq f^*(x)\big)$, where $\mathbb{S}$ denotes the set of permutations of $\{1, \ldots, L\}$. Given a small labeled set of data, it is possible to determine the error minimizing permutation by solving the maximum matching problem.

4 One-Versus-All on the Unit Ball

In this section we consider the implicit structure when there exists a consistent linear one-vs-all classifier over the unit ball $\mathcal{B}^d = \{x \in \mathbb{R}^d : \|x\| \leq 1\}$ and the data distribution p is nearly uniform on the positive region $K = \{x \in$

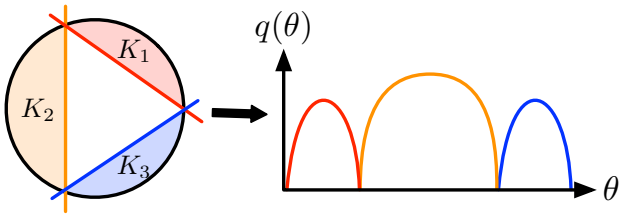

Figure 1: An example 3-class problem on the unit ball with a consistent one-vs-all linear classifier. Also shown is the density q of the projected data when the original density p is uniform on K.

$\mathcal{X} : \exists$ unique h_i s.t. $h_i(x) > 0\}$ for the consistent classifier. Intuitively, these conditions guarantee that w.h.p. an iid sample from p will not contain many nearby pairs of points belonging to different classes because the set of points belonging to any classes i and j can only touch at a point (see Figure 1 for an example in 2 dimensions). When the data distribution is nearly uniform, we can precisely characterize this probabilistic separation in terms of the data density after projecting onto the unit sphere $\mathcal{S}^{d-1}$, which we denote by $q : \mathcal{S}^{d-1} \to [0, \infty)$. The density of the projected data at a point v on the unit sphere depends on the intersection of the ray passing through v with the positive region K. Since the data distribution is both upper and lower bounded, if the ray passes through a thick portion of K then the $q(v)$ must be large and if the ray passes through a thin portion of K, then $q(v)$ must be small.

We use these facts to approximate the set of points belonging to each class by approximating the connected components of the ϵ-level set of q, given by $\{q \geq \epsilon\} = \{v \in \mathcal{S}^{d-1} : q(v) \geq \epsilon\}$ using a threshold clustering algorithm. This algorithm takes a set of points on the unit sphere, discards any points that have too few near neighbors (since these points have low density), and clusters the surviving points. To classify a new example, we output the index of the nearest cluster. Pseudocode is given in Algorithm 1 using $\theta(u, v) = \arccos(u, v)$ to denote the angle between unit vectors u and v and $C(v, r) = \{u \in \mathcal{S}^{d-1} : \theta(v, u) \leq r\}$ to denote the spherical cap of radius r centered at v.

Input: Sample $\mathfrak{X} = \{v_1, \ldots, v_n\}$, connection radius $r_c > 0$, activation radius $r_a > 0$, activation threshold $\tau > 0$.
1. Mark v_i active if $|C(v_i, r_a) \cap \mathfrak{X}|/n \geq \tau$ and inactive otherwise.
2. Construct a graph G on the active points with an edge between v_i and v_j if $\theta(v_i, v_j) < r_c$.
3. Let $\hat{K}_1, \ldots, \hat{K}_N$ be the connected components of G.
4. Output classification rule $\hat{f}(x) = \operatorname{argmin}_{1 \leq i \leq N} \theta(x, \hat{K}_i)$.

Algorithm 1: Robust threshold clustering on the sphere.

More generally, Algorithm 1 can be applied in any metric space $(\mathcal{X}, d)$ by replacing θ with the distance metric d and $C(v, r)$ with the corresponding ball $B(x, r)$. For the robust linkage approach to have low error, each class should have one large connected component in the graph G constructed by the algorithm so that: (1) with high probability a new point in class i will be nearest to that largest component, and (2) the large components of different classes are separated. Intuitively, G will have these properties if each positive region K_i has a connected high-density inner region A_i covering most of its probability mass and when it is rare to observe a point that is close to two or more classes. This notion is formalized below.

Let S be any set in $\mathcal{X}$. We say that a path $\pi : [0, 1] \to \mathcal{X}$ *crosses* S if the path starts and ends in different connected components of the complement of S in $\mathcal{X}$ and we say that the *width* of S is the length of the shortest path that crosses S.

Definition 1. *The sets $A_1, \ldots, A_L$ are (r_c, r_a, τ, γ)-clusterable under probability P if there exists a separating set S of width at least r_c such that: (1) Each A_i is connected; (2) If $x \in \mathcal{X}$ satisfies $d(x, A_i) \leq r_c/3$ then $\mathbb{P}_{X \sim P}(X \in B(x, r_a)) > \tau + \gamma$; (3) If $x \in A_i$ then $\mathbb{P}_{X \sim P}(X \in B(x, r_c/3)) > \gamma$; (4) Every path from A_i to A_j crosses S; and (5) If $x \in S$ then $\mathbb{P}_{X \sim P}(X \in B(x, r_a)) < \tau - \gamma$.*

Note that typically there must be a gap between the set A_i and the set S in order to satisfy the probability requirements (i.e., the set S will be smaller than $\mathcal{X} - \bigcup_{i=1}^{L} A_i$). The first three properties ensure that each class will have one large connected component and the remaining two properties ensure that these connected components will be disconnected. Following an analysis similar to that of the cluster tree algorithm of Chaudhuri and Dasgupta [2010] gives the following result:

Lemma 1. *For any $\epsilon > 0$, suppose each positive region K_i has a subset A_i such that the $\{A_i\}_{i=1}^{L}$ are (r_c, r_a, τ, γ)-clusterable and $P(K - \bigcup_i A_i) < \epsilon$. Then for any $\delta > 0$, let $\hat{f}$ be the output of Algorithm 1 run on a sample of size $n = O(\frac{1}{\gamma^2}(D \ln \frac{D}{\gamma} + \ln \frac{1}{\delta}))$, where D is the VC-dimension of balls in $\mathcal{X}$ ($D = d + 1$ in $\mathbb{R}^d$), with parameters r_c, r_a, and τ. Then with prob. $\geq 1 - \delta$, we have $\mathrm{err}_P(\hat{f}) < \epsilon$.*

Proof. Let $\mathfrak{X}$ denote the sample of size n and $\hat{P}$ denote the empirical measure induced by $\mathfrak{X}$. For each class i, let $\hat{K}_i = \{x \in \mathfrak{X} : d(x, A_i) < r_c/3\}$ be the set of samples within distance $r_c/3$ of A_i for each class $i = 1, \ldots, L$. First, we show that with probability at least $1 - \delta$, the graph G constructed by Algorithm 1 has the following properties:

1. **(Completeness)** Every sample in $\hat{K}_i$ is active and included in the graph G.
2. **(Separation)** For all pairs of classes $i \neq j$, there is no path in G from $\hat{K}_i$ to $\hat{K}_j$.
3. **(Connectedness)** For every class i, the set $\hat{K}_i$ is connected in G.

We use a standard VC bound [Vapnik and Chervonenkis, 1971] to relate the probability constraints in the clusterability definition to the empirical measure $\hat{P}$. For our value of n we have

$$\mathbb{P}\big(\sup_{x,r} \big|\hat{P}(B(x,r)) - P(B(x,r))\big| > \gamma\big) < \delta.$$

This implies that with probability at least $1 - \delta$ for all points x we have: (1) if $d(x, A_i) \leq \frac{r_c}{3}$ for any i then $\hat{P}(B(x, r_a)) > \tau$; (2) if $x \in S$ then $\hat{P}(B(x, r_a)) < \tau$; and (3) if $x \in A_i$ for any i then $\hat{P}(B(x, \frac{r_c}{3})) > 0$. We now use these facts to prove that the graph G has the completeness, separation, and connectedness properties.

Completeness follows from the fact that every sample $x \in \hat{K}_i$ is within distance $r_c/3$ of A_i and therefore $\hat{P}(B(x, r_a)) > \tau$.

To show separation, first observe that every sample $z \in \mathfrak{X}$ that belongs to S will be marked as inactive, since $\hat{P}(B(z, r_a)) < \tau$. Now let $x \in \hat{K}_i$ and $x' \in \hat{K}_j$ for $i \neq j$. Since the graph G does not contain any samples in the set S, any path in G from x to x' must have one edge that crosses S. Since the width of S is at least r_c, this edge would not be included in the graph G, and therefore G does not include a path from x to x'.

To show connectedness, let x and x' be any pair of samples in $\hat{K}_i$ and let v and v' be their nearest points in A_i, respectively. By definition of $\hat{K}_i$, we know that $d(x, v) < r_c/3$ and $d(x', v') < r_c/3$. Since A_i is a connected set, there is a path $\pi : [0, 1] \to A_i$ in A_i starting at v and ending at v'. Cover the path π with a sequence of points $z_1, \ldots, z_k$ such that $d(z_j, z_{j+1}) < r_c/3$ for all j and the path π is covered by the balls $B(z_j, r_c/3)$. Further, choose $z_1 = v$ and $z_k = v'$. Since each point z_j belongs to A_i, the empirical probability mass of the ball $B(z_j, r_c/3)$ is non-zero, which implies that it must contain at least one sample point, say $y_j \in \mathfrak{X}$. We may take $y_1 = x$ and $y_k = x'$. Since every sample $y_1, \ldots, y_k$ is within distance $r_c/3$ of A_i, they are all active and included in the graph G. Moreover, since $d(y_j, y_{j+1}) < r_c$, we have that the path $x = y_1 \to \cdots \to y_k = x'$ is a path connecting x and x' in G, as required.

Finally we show that the output hypothesis $\hat{f}$ has small error. Let x be any new sample drawn from P. Then, with probability at least $1 - \epsilon$ it will belong to one of the sets A_i. In this case, the ball $B(x, r_c/3)$ has non-zero empirical probability mass so it must contain at least one active sample in G, say z. Since z is within distance $r_c/3$ of A_i, it belongs to the $\hat{K}_i$, the large connected component of G corresponding to class i. Now let z^* be the nearest sample to x. We know that $d(x, z^*) \leq d(x, z) < r_c/3$ and therefore, by the triangle inequality, $d(z, z^*) < r_c$ and we have that z and z^* are connected in G. It follows that with probability at least $1 - \epsilon$, the point x will be assigned to the large connected component corresponding to the class of i. Therefore, with probability at least $1 - \epsilon$, the rule $\hat{f}$ will be consistent with f^*. □

In the remainder of this section, we apply Lemma 1 to the case of one-vs-all learning on the unit ball. Formally, we assume that there exist L linear separators $h_i(x) = w_i^\top x - b_i$, $i = 1, \ldots, L$, such that a point $x \in \mathcal{B}^d$ belongs to

class i iff $h_i(x) > 0$. For simplicity, we assume that no class contains the origin, which implies that $b_i > 0$ for all i. Define $K_i = \{x \in \mathcal{B}^d : h_i(x) > 0\}$ to be the set of points that belong to class i and $K = \bigcup_{i=1}^L K_i$. Finally, suppose that the data density p is *nearly uniform* on K; that is, $p(x) = 0$ for all $x \notin K$ and for $x \in K$ we have $p(x) \in [c_{\text{lb}}, c_{\text{ub}}]$ for constants $0 < c_{\text{lb}} \leq c_{\text{ub}}$. In the special case where p is uniform on K, we have $c_{\text{ub}} = c_{\text{lb}} = 1/\operatorname{Vol}(K)$. We let $q : \mathcal{S}^{d-1} \to [0, \infty)$ denote the density of a point sampled from p then projected onto $\mathcal{S}^{d-1}$ (i.e., if $x \sim p$ and $v = x/\|x\|$ then v is distributed according to q). The density q is defined with respect to the uniform distribution on the sphere $\mathcal{S}^{d-1}$, which we denote by $\mu_\circ$. The following result establishes simple upper and lower bounds on the projected density q that will be used to show that the projected density satisfies the condition sof Lemma 1.

Lemma 2. *Let $h_i(x) = w_i^\top x - b_i$ with $b_i > 0$ for $i = 1, \ldots, L$ be linear separators defining a consistent one-vs-all classifier, suppose that p is a nearly uniform density function taking values in $[c_{\text{lb}}, c_{\text{ub}}]$ supported on $\bigcup_i K_i$ where $K_i = \{x \in B : h_i(x) > 0\}$, and let q be the density function after projecting onto the unit sphere $\mathcal{S}^{d-1}$. For every $v \in \mathcal{S}^{d-1}$, if $v \notin \bigcup_i K_i$, then $q(v) = 0$, and if $v \in K_i$, then $q_{\text{lb}}^{(i)}(v) \leq q(v) \leq q_{\text{ub}}^{(i)}(v)$, where $q_{\text{lb}}^{(i)}(v) = c_{\text{lb}} dv_d (1 - (b_i/w_i^\top v)^d)$, $q_{\text{ub}}^{(i)}(v) = \frac{c_{\text{ub}}}{c_{\text{lb}}} q_{\text{lb}}^{(i)}(v)$, and v_d is the volume of the unit ball $\mathcal{B}^d$.*

Proof. Let $X \sim p$ be and set $V = X/\|X\|_2$ so that V is a sample from q. For any set $A \subset \mathcal{S}^{d-1}$, we know that $\mathbb{P}(V \in A) = \mathbb{P}(X \in \operatorname{cone}(A))$, which gives

$$\mathbb{P}(V \in A) = \mathbb{P}(X \in \operatorname{cone}(A)) = \int_{x \in \operatorname{cone}(A)} p(x)\, dx = \int_{v \in A} dv_d \int_{r=0}^{\infty} p(rv) r^{d-1}\, dr\, d\mu_\circ(v),$$

where the last inequality follows by a change of variables x to (r, v) where $r = \|x\|_2$ and $v = x/\|x\|_2$. The term r^{d-1} is the determinant of the Jacobian of the change of variables, and the term dv_d, which is the surface area of the unit sphere $\mathcal{S}^{d-1}$, appears since $\mu_\circ$ is normalized so that $\mu_\circ(\mathcal{S}^{d-1}) = 1$. From this, it follows that the density function q can be written as

$$q(v) = dv_d \int_{r=0}^{\infty} p(rv) r^{d-1}\, dr, \tag{1}$$

since integrating this function over any set A gives the probability that V will land in A. From our assumptions on p, we know that

$$p(rv) \geq \sum_{i=1}^{L} \mathbb{I}\{rv \in K_i\} c_{\text{lb}}.$$

Moreover, we can rewrite the indicator as $\mathbb{I}\{rv \in K_i\} = \mathbb{I}\{\frac{b_i}{w_i^\top v} < r \leq 1\}$. Substituting this into (1) gives

$$\begin{aligned}
q(v) &\geq \sum_{i=1}^{L} c_{\text{lb}} dv_d \int_{r=0}^{\infty} \mathbb{I}\{[\}\right] \frac{b_i}{w_i^\top v} < r \leq 1 r^{d-1}\, dr \\
&= \sum_{i=1}^{L} \mathbb{I}\{v \in K_i\} c_{\text{lb}} dv_d \int_{r=b_i/(w_i^\top v)}^{1} r^{d-1}\, dr \\
&= \sum_{i=1}^{L} \mathbb{I}\{v \in K_i\} c_{\text{lb}} dv_d (1 - b_i/(w_i^\top v)^d) \\
&= \sum_{i=1}^{L} \mathbb{I}\{v \in K_i\} q_{\text{lb}}^{(i)}(v).
\end{aligned}$$

Note that the indicator $\mathbb{I}\{v \in K_i\}$ appears in line 2 because the integral is only non-zero when $b_i/(w_i^\top v) < 1$, which is exactly the condition that $v \in K_i$. The upper bound on q follows by an identical argument using the upper bound on $p(rv)$. □

Since the bounds $q_{\text{lb}}^{(i)}(v)$ and $q_{\text{ub}}^{(i)}(v)$ are decreasing functions of the angle $\theta(w_i, v) = \arccos(w_i^\top v)$, the level set $\{q_{\text{lb}}^{(i)} \geq \lambda\} = \{v \in \mathcal{S}^{d-1} : q_{\text{lb}}^{(i)}(v) \geq \lambda\}$ is a spherical cap centered at w_i of angular radius $\rho_{\text{lb}}^{(i)}(\lambda) = \arccos(b_i(1 -$

$\lambda/(c_{\rm lb} dv_d))^{-1/d}$) (and similarly for $q_{\rm ub}^{(i)}$ with radius $\rho_{\rm ub}^{(i)}(\lambda) = \arccos(b_i(1-\lambda/(c_{\rm ub}dv_d))^{-1/d})$). Moreover, since the upper and lower bounds differ by only a multiplicative factor of $c_{\rm ub}/c_{\rm lb}$, when the level λ is small, the level sets $\{q_{\rm lb}^{(i)} \geq \lambda\}$ and $\{q_{\rm ub}^{(i)} \geq \lambda\}$ are very close. It follows that for small λ, the λ-level set of q is approximately the union of L spherical caps with one cap corresponding to the points belonging to each class. Therefore, if we can estimate the connected components of $\{q \geq \lambda\}$ then we can approximate the set of points belonging to each class (up to a permutation of the labels).

Theorem 1. *Under the same conditions as Lemma 2, define* $\beta = \frac{c_{\rm lb}}{3dv_d c_{\rm ub}^2}\frac{\sqrt{1-b_{\min}^2}}{b_{\min}}$, *where* $b_{\min} = \min_i b_i$. *For any* $\epsilon > 0$ *and* $\delta > 0$, *running Algorithm 1 with parameters* $r_a = \Omega(\epsilon\beta)$, $r_c = 2r_a$, $\tau = \frac{c_{\rm ub}V^d(r_a)}{2c_{\rm ub}}\epsilon$, *and* $\gamma = \frac{c_{\rm ub}V^d(r_c/3)}{4c_{\rm ub}}\epsilon = \Omega\big((\epsilon\beta)^d\big)$, *where* $V^d(r)$ *is the* $\mu_\circ$*-measure of a spherical cap of radius* r, *on a sample of size* $n = O\big(\frac{1}{(\epsilon\beta)^{2d}}\big(d\ln\frac{d}{(\epsilon\beta)^d} + \ln\frac{1}{\delta}\big)\big)$ *drawn from* q *will output a classifier* $\hat{f}$ *with* $\mathrm{err}_P(\hat{f}) \leq \epsilon$ *with probability at least* $1-\delta$.

Proof. To apply Lemma 1, we need to construct sets $A_1, \ldots, A_L$ so that all but at most ϵ of the probability mass of q is contained in the A_i sets and the sets $A_1, \ldots, A_L$ are (r_c, r_a, τ, γ)-clusterable. For each class i, we define $A_i = \{q_{\rm ub}^{(i)} \geq \epsilon\}$ to be the ϵ level-set of our upper bound $q_{\rm ub}^{(i)}$. This is a spherical cap of radius $\rho_{\rm ub}^{(i)}(\epsilon)$ centered at w_i, the weight vector defining the linear separator for class i. By construction, any point v not contained in $\bigcup_i A_i$ has $q(v) \leq \epsilon$, which gives that $\mathbb{P}_{V\sim q}(V \notin \bigcup_i A_i) \leq \epsilon$, so it remains to show that these sets are clusterable.

Define $\tilde{\epsilon} = \frac{c_{\rm lb}}{c_{\rm ub}}\epsilon$ and suppose for the moment that we can find an activation radius $r_a > 0$ small enough so that the following inequalities hold for all classes $i = 1, \ldots, L$:

$$\frac{5}{3}r_a \leq \rho_{\rm lb}^{(i)}\big(\frac{3\tilde{\epsilon}}{4}\big) - \rho_{\rm ub}^{(i)}(\epsilon) \qquad \text{and} \qquad 2r_a \leq \rho_{\rm ub}^{(i)}(0) - \rho_{\rm ub}^{(i)}\big(\frac{\tilde{\epsilon}}{4}\big).$$

Given such an activation radius, we will show that the sets $A_1, \ldots, A_L$ are (r_c, r_a, τ, γ)-clusterable, where $r_c = 2r_a$, $\tau = \frac{\tilde{\epsilon}V^d(r_a)}{2}$, and $\gamma = \frac{\tilde{\epsilon}V^d(r_c/3)}{4}$ and the separating set is $S = \{v \in \mathcal{S}^{d-1} : \theta(v, w_i) \geq \rho_{\rm ub}^{(i)}(0) - r_a \text{ for all } i\}$:

1. *Connection:* Each A_i set is a spherical cap and therefore connected.
2. *High-density near A_i:* Let $v \in \mathcal{S}^{d-1}$ be such that $\theta(v, A_i) < r_c/3$ and let $u \in C(v, r_a)$ be any point in the spherical cap of angular radius r_a about v. By the triangle inequality, we know that $\theta(w_i, u) \leq \theta(w_i, v) + \theta(v, u) \leq \rho_{\rm ub}^{(i)}(\epsilon) + \frac{5}{3}r_a \leq \rho_{\rm lb}^{(i)}(\frac{3\tilde{\epsilon}}{4})$. This implies that $q(u) \geq \frac{3\tilde{\epsilon}}{4}$ for all points in $C(v, r_a)$ and therefore $\mathbb{P}_{V\sim q}(V \in C(v, r_a)) \geq \frac{4\tilde{\epsilon}}{3}V^d(r_a) \geq \tau + \gamma$.
3. *High-density inside A_i:* Now let $v \in A_i$. Since $r_c/3 < r_a$, the above arguments show that $q(u) \geq \frac{4\tilde{\epsilon}}{3}$ for all points $u \in C(v, r_c/3)$ and therefore $\mathbb{P}_{V\sim q}(V \in C(v, r_c/3)) \geq \frac{4\tilde{\epsilon}}{3}V^d(r_c/3) \geq \gamma$.
4. *Separation by the set S:* For each class i, the set S contains the annulus $\{v \in \mathcal{S}^{d-1} : \rho_{\rm ub}^{(i)}(0) - r_a \leq \theta(w_i, v) \leq \rho_{\rm ub}^{(i)}(0)\}$ which has width r_a. Any path from one A_i to another A_j must cross two such annuli, each of width r_a, so the length of the path crossing S is at least $2r_a = r_c$.
5. *Low density inside S:* Finally, let v be any point in the set S and let $u \in C(v, r_a)$. For any class i, the reverse triangle inequality gives that $\theta(w_i, v) \geq \theta(v, w_i) - \theta(u, w_i) \geq \rho_{\rm ub}^{(i)}(0) - 2r_a \geq \rho_{\rm ub}^{(i)}(\frac{\tilde{\epsilon}}{4})$. Since this is true for all classes i, we have $q(v) \leq \frac{\tilde{\epsilon}}{4}$ and therefore $\mathbb{P}_{V\sim q}(V \in C(v, r_a)) \leq \frac{\tilde{\epsilon}}{4}V^d(r_a) \leq \tau - \gamma$.

It follows that the sets $A_1, \ldots, A_L$ are (r_c, r_a, τ, γ)-clusterable and it only remains to find an activation radius r_a that satisfies the above inequalities. Since the robust linkage algorithm needs to estimate the probability mass of balls to within error $\gamma = \frac{\tilde{\epsilon}V^d(2r_a/3)}{4}$, we want this activation radius to be not too small.

Taking the first order Taylor expansion of the $\rho_{\rm lb}^{(i)}$ and $\rho_{\rm ub}^{(i)}$ functions, we have:

$$\rho_{\rm lb}^{(i)}(\lambda) = \arccos(b_i) - \frac{b_i}{\sqrt{1-b_i^2}}\frac{1}{c_{\rm lb}dv_d}\lambda + O(\lambda^2)$$
$$\rho_{\rm ub}^{(i)}(\lambda) = \arccos(b_i) - \frac{b_i}{\sqrt{1-b_i^2}}\frac{1}{c_{\rm ub}dv_d}\lambda + O(\lambda^2),$$

as $\lambda \to 0$. Therefore, we have that

$$\rho_{\mathrm{lb}}^{(i)}(3\tilde{\epsilon}/4) - \rho_{\mathrm{ub}}^{(i)}(\epsilon) = \frac{1}{4dv_d c_{\mathrm{ub}}} \cdot \frac{b_i}{\sqrt{1-b_i^2}}\epsilon + O(\epsilon^2)$$

and

$$\rho_{\mathrm{ub}}^{(i)}(0) - \rho_{\mathrm{ub}}^{(i)}(\tilde{\epsilon}/4) = \frac{c_{\mathrm{lb}}}{4dv_d c_{\mathrm{ub}}^2} \cdot \frac{b_i}{\sqrt{1-b_i^2}}\epsilon + O(\epsilon^2),$$

which shows that it is sufficient to set $r_a = \frac{3c_{\mathrm{lb}}}{20dv_d c_{\mathrm{ub}}^2} \cdot \frac{b_i}{\sqrt{1-b_i^2}}\epsilon + O(\epsilon^2) = \Omega(\frac{c_{\mathrm{lb}}}{dv_d c_{\mathrm{ub}}^2} \cdot \frac{b_i}{\sqrt{1-b_i^2}}\epsilon) = \Omega(\epsilon\beta)$ as $\epsilon \to 0$. The sample complexity bound follows from the fact that $V^d(\alpha) = \Omega((\alpha/2)^{(d-1)})$ as $\alpha \to 0$.

□

5 Error Correcting Output Codes

Next we consider the implicit structure when there exists a consistent *error correcting* output code classifier with code matrix $C \in \{\pm 1\}^{L \times m}$ and linear functions $h_1, \ldots, h_m$. We assume that for any point x, the consistent output code classifier makes at most β mistakes when predicting the codeword for x (that is, the Hamming distance between $h(x)$ and C_y is at most β, where y is the label of x), and we assume that the Hamming distance between the code words is at least $2\beta + d + 1$. When $\beta > 0$, the linear separators $h_1, \ldots, h_m$ are permitted to make mistakes when predicting the semantic feature values for a new example. This condition captures the intuition that output codes with well separated code words can correct a small number of errors.

For any points x and x' that belong to different classes, we must have that the Hamming distance between $h(x)$ and $h(x')$ is at least $d + 1$. Under the extremely mild condition that the linear separators $h_1, \ldots, h_m$ defining the consistent classifier are in general position (i.e., at most d intersect at any given point), this implies that there is a minimum distance $g > 0$ such that points belonging to different classes must be at least distance g apart. We exploit this structure under various assumptions on the data density p.

Input: Sample $\mathfrak{X} = \{x_1, \ldots, x_n\}$, connection radius $r_c > 0$.
1. Construct a graph G on the samples with an edge between x_i and x_j if $\|x_i - x_j\| < r_c$.
2. Let $\hat{K}_1, \ldots, \hat{K}_N$ be the connected components of G.
3. Output $\hat{f}(x) = \mathrm{argmin}_{1 \le i \le N}\, d(x, \hat{K}_i)$, where $d(x, S) = \min_{x' \in S} \|x - x'\|$.

Algorithm 2: Threshold clustering algorithm.

Nearly uniform distribution on round sets: When the data distribution is nearly uniform, show that a simple threshold clustering algorithm (pseudocode given in Algorithm 2) can recover the classes by clustering an unlabeled sample. We assume that the data distribution is nearly uniform on the set $K = \bigcup_{i=1}^{L} K_i$ where $K_i = \{x \in \mathcal{X} : d_{\mathrm{Ham}}(h(x), C_i) \le \beta\}$ is the set of points belonging to class i. To get finite sample guarantees, we borrow the notion of (ϵ, σ)-roundness from Blum and Chawla [2001]: For any $\sigma > 0$, the σ-interior of a set A is the set $\mathrm{int}_\sigma(A) = \{x \in A : \exists y \in A \text{ s.t. } x \in B(y, \sigma) \subset A\}$ and the σ-tendrils of A are the points further than distance σ from the σ interior: $\mathrm{tendril}_\sigma(A) = \{x \in A : d(x, \mathrm{int}_\sigma(A)) > \sigma\}$.

Definition 2. *A set A is (σ, ϵ)-round with respect to probability P if $\mathrm{int}_\sigma(A)$ is non-empty, connected, and $\mathbb{P}_{x \sim P}(x \in \mathrm{tendril}_\sigma(A)) \le \epsilon$.*

Intuitively, the set A is (ϵ, σ)-round if all but an ϵ-mass of its points can be reached by a ball of radius σ contained in A. Following a similar proof technique as Blum and Chawla, we obtain the following result.

Theorem 2. *Suppose there exists a linear output code classifier that makes at most β errors predicting the code word of a new example, the Hamming distance between code words is at least $2\beta + d + 1$, and the data density p is nearly*

uniform over the set $K = \bigcup_i K_i$ where $K_i = \{x \in \mathcal{X} : d_{\mathrm{Ham}}(h(x), C_i) \leq \beta\}$ taking values in $[c_{\mathrm{lb}}, c_{\mathrm{ub}}]$. Let $g > 0$ be the gap between classes. For any $\epsilon > 0$, suppose there exists a number $0 < \sigma_\epsilon < g/8$ such that the positive regions $K_1, \ldots, K_L$ are $(\epsilon, \sigma_\epsilon)$-round. For any $\delta > 0$, let $\hat{f}$ be the output of Algorithm 2 run with $r_c = g/2$ on a sample of size $n \geq \frac{1}{c_{\mathrm{lb}} \sigma_\epsilon^d v_d} \ln \frac{2^d \operatorname{Vol}(K)}{\delta \sigma_\epsilon^d v_d}$, where v_d is the volume of the unit ball $\mathcal{B}^d$. Then with probability at least $1 - \delta$, we have $\operatorname{err}_P(\hat{f}) \leq L\epsilon$.

Proof. We begin by covering the σ_ϵ-interior of the positive region K with finite set of balls of radius σ_ϵ. Consider the following greedy procedure: while there exists a point x in the σ_ϵ-interior of K that is not covered by our current set of balls, add the ball $B(x, \sigma_\epsilon)$ to the cover. This procedure only terminates once the σ_ϵ-interior of K is covered, so it remains to bound the number of iterations before it terminates. Suppose that after t iterations we have added balls with centers $c_1, \ldots, c_t$. Since each point c_j is in the σ_ϵ-interior of K, we know that $B(x_j, \sigma_\epsilon) \subset K$. Moreover, since no ball of radius σ_ϵ contains the center of any other ball, it follows that the balls with centers $c_1, \ldots, c_t$ of radius $\sigma_\epsilon/2$ are disjoint (and contained in K). It follows that we have

$$\operatorname{Vol}(K) \geq \operatorname{Vol}\Big(\bigcup_{j=1}^{t} B(c_j, \sigma_\epsilon/2) \Big) = \sum_{j=1}^{t} \operatorname{Vol}\big(B(c_j, \sigma_\epsilon/2)\big) = t(\sigma_\epsilon/2)^d v_d.$$

Rearranging this inequality gives an upper bound on the number of iterations before terminating:

$$t \leq \frac{2^d \operatorname{Vol}(K)}{\sigma_\epsilon^d v_d}.$$

For any ball B of radius σ_ϵ contained in the positive region K, we have $\mathbb{P}_{x \sim p}(x \in B) \geq c_{\mathrm{lb}} \sigma_\epsilon^d v_d$. Therefore, the probability that an iid sample of size n from p does not contain any sample in B is at most $(1 - c_{\mathrm{lb}} \sigma_\epsilon^d v_d)^n \leq \exp(-n c_{\mathrm{lb}} \sigma_\epsilon^d v_d)$. Taking the union bound over the balls in our cover, the probability that there exists any ball in the cover that does not contain a sample is at most $\frac{2^d \operatorname{Vol}(K)}{\sigma_\epsilon^d v_d} \exp(-n c_{\mathrm{lb}} \sigma_\epsilon^d v_d)$. Setting n as in the Theorem statement makes this upper bound equal to δ. Assume that this high probability event holds for the remainder of the proof.

Next we show that the graph G constructed by Algorithm 1 has one large connected component for each of the classes and that these components are disconnected from each other. Fix any class i and let $B_1, \ldots, B_k$ be the balls from our covering that cover the σ_ϵ-interior of K_i and let $\hat{K}_i = \mathfrak{X} \cap (\bigcup_{j=1}^k B_j)$ be the set of sample points that land in those balls. We will show that all samples in $\hat{K}_i$ are connected in G. Let $x, y \in \hat{K}_i$. The set $\bigcup_{j=1}^k B_j$ is connected, since the σ_ϵ-interior of K_i is connected, and it follows that there is a path $\pi : [0, 1] \to \bigcup_{j=1}^k B_j$ from x to y that stays in our covering of the σ_ϵ-interior of K_i. The path π must move through some sequence of balls $B_{j_1}, \ldots, B_{j_N}$ (possibly with repetition). By the high probability event, each ball contains at least one sample point. Let $v_1, \ldots, v_N$ be any sequence of samples with $v_l \in B_{j_l}$. Without loss of generality, we may assume that $v_1 = x$ and $v_N = y$. Since B_{j_l} and $B_{j_{l+1}}$ touch, the triangle inequality guarantees that $d(v_j, v_{j+1}) \leq 4\sigma_\epsilon < r_c$, which shows that the path $x = v_1 \to v_2 \cdots \to v_{N-1} \to v_N = y$ is a path in G from x to y. Finally, since the distance between any pair of positive regions is at least $g > r_c$, the graph G will never conntain a path from one $\hat{K}_i$ to another $\hat{K}_j$.

Now suppose that x is any new point that belongs to K_i but not the σ_ϵ-tendrils of K_i. We will show that the nearest sample in $\mathfrak{X}$ to x is connected to $\hat{K}_i$. By definition of the σ_ϵ-tendrils, the distance from x to the σ_ϵ-interior of K_i is at most σ_ϵ, so we can find some point $y \in \operatorname{int}_{\sigma_\epsilon}(K_i)$ so that $d(x, y) < \sigma_\epsilon$. The point y must belong to some ball B_j from our covering of the σ_ϵ-interior of K_i, and the ball B_j contains at least one sample, say z. By the triangle inequality, we have that $d(x, z) \leq 2\sigma_\epsilon$. Now let z^* be the closest sample in $\mathfrak{X}$ to x. We know that $d(x, z^*) \leq d(x, z) = 2\sigma_\epsilon$, from which it follows that $d(z, z^*) \leq 4\sigma_\epsilon < r_c$, and z^* is connected to z in the graph G. Since $z \in \hat{K}_i$, it follows that z^* is connected to $\hat{K}_i$.

The above arguments show that the classification rule $\hat{f}$ will assign different labels to any points that belong to different classes, and it will assign the same label to any point in K_i, except possibly for those that belong to the σ_ϵ-tendrils. Since the total probability mass of the σ_ϵ-tendrils is $L\epsilon$, it follows that the error of $\hat{f}$ is at most $L\epsilon$. □

Thick level sets: Next, we significantly generalize the above result by removing the assumption that the data distribution is uniform over the positive region K and designing an algorithm for which the gap g does not need to be

known in advance. Instead, we will assume only that the data distribution has a density $p : \mathcal{X} \to [0, \infty)$ that satisfies a thick level set condition. Intuitively, this condition requires that the level sets $\{p \geq \lambda\}$ do not have bridges or cusps that are too thin.

The added challenge in this setting is that the points belonging to class i in an iid sample from p might not form a single connected cluster. The basic idea of our algorithm is to use a small amount of labeled data to find which clusters belong to which classes. Specifically, our algorithm computes the single linkage hierarchical clustering of an unlabeled dataset, queries the labels of a small number of examples, and finds the smallest pruning of the hierarchical clustering where each cluster is label-homogeneous. Recall that the single linkage hierarchical clustering is a binary tree where each node corresponds to a subset of data with the root being the entire dataset and the leaves being individual examples. The tree is built from the leaves up by repeatedly joining the two nodes having the closest pair of points. A pruning of the hierarchical clustering is a collection of nodes in the tree that partition the data. Our algorithm finds the smallest pruning (in terms of number of clusters) for which each cluster contains labeled examples from exactly one class. A new example x is classified by assigning it to the label of the nearest cluster. See Algorithm 3 for pseudocode.

Input: Sample $\mathfrak{X} = \{x_1, \ldots, x_n\}$, estimate $\hat{N}$ of number of clusters, $\epsilon > 0$, $\delta > 0$.

1. Let T be the single linkage hierarchical clustering of $\mathfrak{X}$.
2. Query the labels of a random subset of $\mathfrak{X}$ of size $\frac{1}{2\epsilon} \ln \frac{2\hat{N}}{\delta}$ (without replacement).
3. Let $\hat{K}_1, \ldots, \hat{K}_{\hat{N}}$ be the clusters obtained by finding the smallest pruning of T for which each cluster contains labeled examples from at most one class.
4. Output classification rule $f(x)$ that assigns x to the same class as the nearest labeled cluster.

Algorithm 3: Active Hierarchical Single Linkage

Our analysis of this algorithm has two main steps. First, we show that for a well chosen connection radius $r_c < g$, running the threshold clustering algorithm (Algorithm 2) on an iid sample would output a clustering where every cluster is label homogeneous and the number of clusters is approximately equal to the number of connected components of $\{p \geq \epsilon\}$. Next, as long as we query the label from every large cluster that would be output by Algorithm 2, our pruning of the hierarchical single linkage clustering will be label homogeneous and have at most the same number of clusters. By querying the labels of a large enough set of examples, we can ensure that this happens with high probability. Both of these results use the following thick level set condition borrowed from Steinwart [2015].

Definition 3. *A density function p has thick level sets up to level $\lambda_0 \geq 0$ if there exist constants $C > 0$ and $\sigma_0 > 0$ such that for all levels $\lambda \in [0, \lambda_0]$ and all $\sigma \in (0, \sigma_0]$, we have: (1) the σ-interior of $\{p \geq \lambda\}$ is non-empty and (2)* $\sup_{x \in \{p \geq \lambda\}} d\big(x, \mathrm{int}_\sigma(\{p \geq \lambda\})\big) \leq C\sigma$.

The parameter C measures how pointy the level sets of p can be. For example, suppose that the λ-level set of p is a square in $\mathbb{R}^2$. Then the σ-interior of the level set will be another square whose side length is reduced by 2σ and the furthest point in $\{p \geq \lambda\}$ from the σ-interior is $\sqrt{2}\sigma$, so in this case, $C = 2$. Analogously, if the level set is a hypercube in d-dimensions, then $C = \sqrt{d}$. On the other hand, if the level set is a ball in $\mathbb{R}^d$, then the thickness constant C is 1. This matches our intuition that the square is pointier than the unit ball.

The goal of both Definition 2 and Definition 3 is to ensure that the points further than distance σ from the σ-interior of the level sets of p will not cause problems for our clustering algorithms. The roundness condition assumes that those points have small probability mass, while the thickness condition assumes that they are not too far from the σ-interior (at most distance $C\sigma$). In the case when $C = 1$, the thickness condition implies that each connected component of the level set will be $(\sigma, 0)$-round, but for $C > 1$ the two conditions are incomparable. Theorem 5 in Section 8 of the supplementary material proves an analogous result to Theorem 2 using the thickness condition in place of roundness when the distribution is uniform over K.

Our first result does not depend on the thick level set condition, but shows that we can cover the σ-interior of the λ-level set of p with a relatively small number of balls of radius σ. Moreover, if we see a sufficiently large sample drawn from p, then every ball in this covering will contain at least one sample point.

Lemma 3. *For any $\lambda > 0$ and any $\sigma > 0$, there is a covering of the σ-interior of $\{p \geq \lambda\}$ by balls of radius σ of size at most $N = 2^d/(\lambda\sigma^d v_d)$. Moreover, for any $\delta > 0$, if we see an iid sample $x_1, \ldots, x_n$ drawn from p of size $n \geq \frac{1}{\lambda\sigma^d v_d}\big(\ln \frac{2^d}{\lambda\sigma^d v_d} + \ln \frac{1}{\delta}\big)$, then with probability at least $1 - \delta$, every ball in our cover will contain at least one sample.*

Proof. Consider the following greedy procedure: as long as there exists any point x in the σ-interior of $\{p \geq \lambda\}$ that is not covered by our current cover, add the ball $B(x, \sigma)$ to the cover. Suppose after t additions, we have added balls with centers $c_1, \ldots, c_t$. Since no ball in the cover contains the center of any other, we know that the distance between any pair of centers is at least $\sigma/2$. Therefore, the balls $B(c_1, \sigma/2), \ldots, B(c_t, \sigma/2)$ are disjoint and contained in $\{p \geq \lambda\}$ (since their centers are in the σ-interior of $\{p \geq \lambda\}$). This gives

$$1 \geq P\bigg(\bigcup_{i=1}^{t} B(c_i, \sigma/2)\bigg) = \sum_{i=1}^{t} P\big(B(c_i, \sigma/2)\big) \geq t\lambda(\sigma/2)^d v_d.$$

Rearranging the above inequality gives an upper bound of $2^d/(\lambda\sigma^d v_d)$ on the number of centers added into our cover.

For any center c_i in our cover, the ball $B(c_i, \sigma)$ is contained in $\{p \geq \lambda\}$, which implies that its probability mass is at least $\lambda\sigma^d v_d$. Therefore the probability that none of the n iid samples falls in $B(c_i, \sigma)$ is at most $(1 - \lambda\sigma^d v_d)^n \leq \exp(-n\lambda\sigma^d v_d)$. Applying the union bound over the balls in our cover, we have

$$\mathbb{P}\big(\exists i \text{ s.t. } B(c_i, \sigma) \text{ does not contain a sample}\big) \leq \frac{2^d}{\lambda\sigma^d v_d} \exp(-n\lambda\sigma^d v_d).$$

Substituting the provided lower bound on n results in the right hand side being equal to δ. □

Next we combine Lemma 3 together with the thick level set condition to show that running single linkage on a large enough sample with a large enough connection radius will not subdivide the connected components of $\mathrm{int}_\sigma(\{p \geq \lambda\})$.

Lemma 4. *For any $\lambda \in (0, \lambda_0]$, $\sigma \in (0, \sigma_0]$, and $\delta > 0$, let $\mathfrak{X}$ be an iid sample drawn from p of size*

$$n \geq \frac{1}{\lambda\sigma^d v_d}\left(\ln \frac{2^d}{\lambda\sigma^d v_d} + \ln \frac{1}{\delta}\right)$$

and let G be the graph whose vertices are the samples $\mathfrak{X}$ with an edge between samples x and x' if $\|x - x'\| \leq r_c$, where the connection radius is $r_c = (2C + 6)\sigma$. Let $A_1, \ldots, A_N$ be the connected components of the σ-interior of $\{p \geq \lambda\}$ and define $C_i = \{x \in \{p \geq \lambda\} : d(x, A_i) \leq C\sigma\}$ for $i = 1, \ldots, N$. Then with probability at least $1 - \delta$, for every i, the set of samples

$$\hat{C}_i = \{x \in \mathfrak{X} : x \text{ is connected to a point in } C_i \text{ in } G\}$$

will be a connected component of the graph G. Moreover, for any new point $x \in C_i$, the closest sample in the graph G will belong to $\hat{C}_i$.

Proof. By Lemma 3, we know there exists a covering of the σ-interior of $\{p \geq \lambda\}$ by balls of radius σ and, with probability at least $1 - \delta$, each ball in the cover will contain at least one sample. Assume that this high probability event holds.

Consider any connected component A_i of the σ-interior of $\{p \geq \lambda\}$ and let x, x' be any pair of samples that belong to C_i. We will show that x and x' must be connected in the graph G. By definition of C_i, there exist points z and z' in A_i such that $\|x - z\| \leq C\sigma$ and $\|x' - z'\| \leq C\sigma$. Since both z and z' belong to the connected set A_i, there is a path $\pi : [0, 1] \to A_i$ starting at z and ending at z'. And, since the set A_i is contained in the σ-interior of $\{p \geq \lambda\}$, our covering of $\mathrm{int}_\sigma(\{p \geq \lambda\})$ is also a covering of A_i; therefore the path π starts at z and passes from one ball in our cover to another until it reaches the point z'. Since each ball contains at least one sample, we can find a sequence of samples $z_1, \ldots, z_k \in S$ such that $\|z_i - z_{i+1}\| \leq 4\sigma$ for all i, $\|x - z_1\| \leq (C + 2)\sigma$, and $\|z_k - z'\| \leq (C + 2)\sigma$. The furthest distance between any consecutive pair of points along the path $x \to z_1 \to \cdots \to z_k \to x'$ is $\max\big((C + 2)\sigma, 4\sigma\big) \leq (C + 6)\sigma \leq r_c$, and therefore this path exists in the graph G and connects x to x'.

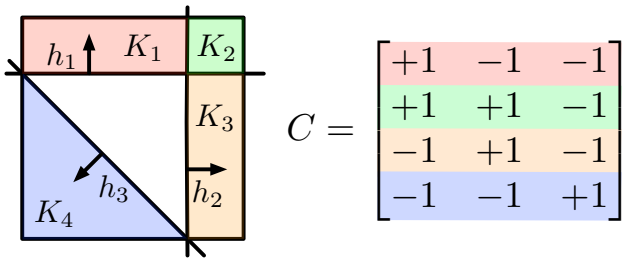

Figure 2: An example of the boundary features problem. The arrows indicate the positive side of the linear functions.

Finally, let x be any new point in C_i. By definition of C_i, there exists a point z in A_i with $\|x - z\| \leq C\sigma$. The point z must be contained in some ball from our covering, which must contain at least one sample. Therefore, there exists a sample $\hat{z} \in C_i$ that is within distance at most $(C+2)\sigma$ from the point x. Now let $z^* \in S$ be the nearest point to x. We must have $\|x - z^*\| \leq \|x - \hat{z}\|$ and therefore $\|\hat{z} - z^*\| \leq (2C+4)\sigma < r_c$ and the points z^* and $\hat{z}$ are connected in the graph G. Since $\hat{z}$ was within distance σ from some connected component of the σ-interior of $\{p \geq \lambda\}$, the claim follows. □

Our main result for the thick level set setting is given below.

Theorem 3. *Suppose there exists a linear output code classifier that makes at most β errors predicting the code word of a new example, the Hamming distance between code words is at least $2\beta + d + 1$, and the data density p has thick level sets up to level λ_0 with parameters σ_0 and C. For any $\epsilon > 0$ and $\delta > 0$, let $\lambda = \min\left(\frac{\epsilon}{\mathrm{Vol}(K)}, \lambda_0\right)$ and $\sigma = \min\left(\frac{g}{(2C+6)}, \sigma_0\right)$. Suppose that $\mathrm{int}_\sigma(\{p \geq \lambda\})$ has N connected components, each with probability mass at least ϵ and let $\hat{f}$ be the output of Algorithm 3 run with parameter $\hat{N} \geq N$ on an unlabeled dataset of size $n = O\left(\frac{1}{\epsilon^2 \sigma^d v_d}\left(\ln \frac{2^d}{\epsilon \sigma^d v_d} + \ln \frac{N}{\delta}\right)\right)$. Then with probability $\geq 1 - \delta$, $\hat{f}$ will have error at most ϵ.*

Proof. For any connection radius $r_c > 0$, we can obtain a pruning of the single linkage hierarchical clustering by only merging clusters that are closer than distance r_c. This is equivalent to running the threshold clustering algorithm with connection radius r_c.

First, we know that the pruning at distance $r_c = g$ (the gap distance between classes) will have clusters that contain points from exactly one class, since otherwise this would contradict the fact that points belonging to different classes must be at least distance g apart. Since Algorithm 3 outputs the coarsest clustering where each cluster has exactly one label, it will output a coarsification of the distance-g pruning. Moreover, by our choice of σ, we know that $g \geq (2C+6)\sigma$.

Next, for our choice of n, Lemma 4 guarantees that with probability at least $1 - \delta/2$, for each index $i = 1, \ldots, N$, the set $S \cap C_i$ (for C_i defined in Lemma 4) will already be connected in the distance-$(2C+6)\sigma$ pruning. Our set of labeled examples is an iid sample drawn from P of size $m \geq \frac{2}{\epsilon} \ln \frac{2N}{\delta}$ and, since each C_i set has mass at least ϵ, applying Hoeffding's inequality and the union bound guarantees that with probability at least $1 - \delta/2$, we will see at least one labeled example from each set C_i. Therefore, with probability at least $1 - \delta$, for each i, the set $C_i \cap S$ will be contained in one cluster in the output from Algorithm 3 and the label of that cluster must agree with the label of the points in C_i. It follows that the resulting classification rule will never make mistakes when predicting on points belonging to $\bigcup_{i=1}^N C_i$. Finally, since $\{p \geq \lambda\} \subset \bigcup_{i=1}^N C_i$ and our choice of λ ensures that a sample from p will land in $\{p \geq \lambda\}$ with probability at least $1 - \epsilon$, this concludes the proof. □

Notice that Algorithm 3 queries the labels of $O(\frac{\log 1/\delta}{\epsilon})$ points. The dependence on ϵ here is the same as one would expect from standard supervised learning, since we are in the realizable setting. The main advantage is that the number of label queries is independent of the dimension and other problem parameters.

6 The Boundary Features Condition

In this section we introduce a natural condition on the code matrix called the boundary features condition that captures the intuition that every semantic feature in an output code should be important. This condition does *not* require that the code words are well separated. We assume that there exists a consistent linear output code classifier with code matrix $C \in \{\pm 1\}^{L \times m}$ and linear functions $h_1, \ldots, h_m$ that correctly predicts the code word for new examples (i.e., $\beta = 0$ in the notation of the previous section). Note that such an assumption is necessary if we allow the Hamming distance between code words to be 1, since making a single mistake when predicting the code word for a new example can result in an incorrect classification.

The boundary features condition requires that every semantic feature constrains the support of the data. The set of points that belong to class i is given by $K_i = \{x \in \mathcal{X} : h(x) = C_i\}$, which is a convex polytope whose faces are specified by a subset of the linear separators $h_1, \ldots, h_m$. Geometrically, the boundary features condition requires that for every linear separator h_j, there exists a class i such that h_j defines a face of K_i that is not shared by the polytope for any other class. Moreover, for a scale parameter $R > 0$, there must be a point x on that face such that the ball $B(x, R)$ does not intersect with any other hyperplane (note this condition can always be satisfied for small enough R). See Figure 2 for an example in 2 dimensions. We can also describe the boundary features condition in terms of the code matrix C: for every column j, there exists a row i such that negating C_{ij} results in a code word $c' \in \{\pm 1\}^m$ that is not already a row of C and for which the set $\{x \in \mathcal{X} : h(x) = c'\}$ has non-zero volume. To simplify our algorithm and analysis, we also assume that for any pair of points $x, x' \in \mathcal{X} - K$ with $h(x) \neq h(x')$ must have $\|x - x'\| \geq R$. Intuitively, this guarantees that the set K does not have any corners with acute interior angles. We consider the case when the data distribution is nearly uniform on the set K.

Notice that the above conditions do not require that the positive regions $K_1, \ldots, K_L$ are separated. In fact, if two code words differ on only a single entry, then the distance between them is zero. For example, in Figure 2, a clustering algorithm will not be able to distinguish between classes 1, 2, and 3, since they form a single well-connected region. Therefore, in general it will not be possible to learn in this setting by clustering an unlabeled sample. Instead, we will use the *absence* of data to estimate the linear separators $h_1, \ldots, h_m$ by searching for linear separators which locally have samples on one side but not on the other. Such hyperplanes must be faces of one of the K_i sets, and are therefore candidate approximations for one of the separators $h_1, \ldots, h_m$.

Our hyperplane detection algorithm works by searching for balls of radius r whose centers are sample points (and therefore must belong to K) such that one half of the ball contains very few samples, say less than a τ fraction of the sample set. After applying a uniform convergence argument for half-balls, if a half-ball contains very few sample points then it must be mostly disjoint from the set K. But since its center point belongs to the set K, this means that the hyperplane defining the half-ball must run nearly parallel to the boundary of K, and therefore is a good approximation to at least one of the true hyperplanes. See Figure 3 for examples of half-balls that would pass and fail this test. Given the collection H of hyperplanes produced in this way, the algorithm classifies a new point according to its code word under the hyperplanes in H. More specifically, the set of hyperplanes in H partition the instance space $\mathcal{X}$ into a collection of cells and we assign a different class label to each cell. Pseudocode is given in Algorithm 4 using the following notation: for any center $x \in \mathcal{X}$, radius $r \geq 0$, and direction $w \in \mathcal{S}^{d-1}$, let $B^{1/2}(x, r, w) = \{y \in B(x, r) : w^\top (y - x) > 0\}$ denote the half-ball with center x, radius r, and direction w and define $p^{1/2}(r) = \frac{1}{2} c_{\mathrm{lb}} r^d v_d$ to be a lower bound on the probability mass of a half-ball of radius r contained entirely in the positive region K (even if no such half-ball exists).

Each candidate hyperplane produced by Algorithm 4 is associated with a half-ball that caused it to be included in H. In fact, we can think of the pairs $(\hat{x}, \hat{w})$ in H as either encoding the linear function $\hat{h}(x) = w^\top (x - \hat{x})$ or the half-ball $B^{1/2}(\hat{x}, r, \hat{w})$, where r is the scale parameter of the algorithm. Most of our arguments will deal with the half-balls directly, so we adopt the second interpretation. The analysis of Algorithm 4 has two main steps. First, we show that the face of every half-ball in the set H is a good approximation to at least one of the true hyperplanes, and that every true hyperplane is well approximated by the face of at least one half-ball in H. Second, using the fact that the half-balls in H are good approximations to the true hyperplanes, we argue that the output classifier will only be inconsistent with the true classification rule in a small margin around each of the true linear separators. Then the error of the classification rule is easily bounded by bounding the probability mass of these margins.

To measure the approximation quality, we say that the half-ball $B^{1/2} = B^{1/2}(\hat{x}, r, \hat{w})$ is an α-approximation to the linear function h if $\mathbb{P}_{X \sim B^{1/2}}(\mathrm{sign}(h(X)) = \mathrm{sign}(h(\hat{x}))) \leq \alpha$, where $\mathbb{P}_{X \sim B^{1/2}}$ denotes the probability when X is

Input: Sample $\mathfrak{X} = \{x_1, \ldots, x_n\}$, threshold $\tau > 0$, scale $r > 0$.

1. Initialize set of candidate hyperplanes $H = \emptyset$.
2. For all samples $\hat{x} \in \mathfrak{X}$ with $B(\hat{x}, r) \subset \mathcal{X}$:
   (a) Let $\hat{w} = \operatorname{argmin}_{w \in \mathcal{S}^{d-1}} |B^{1/2}(\hat{x}, r, w) \cap \mathfrak{X}|$.
   (b) If $|B^{1/2}(\hat{x}, r, \hat{w}) \cap \mathfrak{X}|/n < \tau$ then add $(\hat{x}, \hat{w})$ to H.
3. Let $H = \{(\hat{x}_1, \hat{w}_1), \ldots, (\hat{x}_M, \hat{w}_M)\}$ and define $\hat{h}_i(x) = \hat{w}^\top(x - \hat{x})$ for $i = 1, \ldots, M$.
4. Output rule $\hat{f}$ that classifies x according to the codeword $\langle \operatorname{sign}(\hat{h}_1(x)), \ldots, \operatorname{sign}(\hat{h}_M(x)) \rangle$.

Algorithm 4: Plane-detection algorithm.

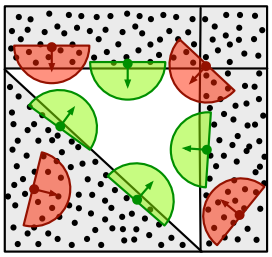

Figure 3: Examples of half-balls that would be included (green) or excluded (red) by the plane detection algorithm.

sampled uniformly from the half-ball $B^{1/2}$. The motivation for this definition is as follows: given any point $\hat{x} \in \mathcal{X}$, the half-ball $B^{1/2}(\hat{x}, r, \hat{w})$ will be an α-approximation to h_i only if $\hat{x}$ is on one side of the decision surface of h_i and all but an α-fraction of the half-ball's volume is on the other side. Intuitively, this means that the face of the half-ball must approximate the decision surface of the function h_i.

The following Lemma shows that when Algorithm 4 is run with appropriate parameters and on a large enough sample drawn from the data distribution, then with high probability the algorithm will include at least one half-ball in H α-approximating each true hyperplane h_i and every half-ball in H will be an α-approximation to at least one true hyperplane. Recall that $p^{1/2}(r) = \frac{1}{2} c_{\text{lb}} r^d v_d$ is a lower bound on the probability mass of a half-ball of radius r contained in the set K.

Lemma 5. *Fix any $\alpha > 0$ and confidence parameter $\delta > 0$. Let H be the set of half-balls produced by Algorithm 4 when run with parameters $r = R/2$ and $\tau = \frac{1}{2}\alpha p^{1/2}(r)$ on a sample of size $n = O(\frac{1}{\gamma^2}(\ln^2 \frac{d}{\gamma} + \ln \frac{1}{\delta}))$ where $\gamma = \frac{2}{5}\tau = \frac{1}{5}\alpha p^{1/2}(r)$. Then with probability at least $1 - \delta$, every half-ball in H will be an α-approximation to at least one true hyperplane h_i, and every true hyperplane h_i will be α-approximated by at least one half-ball in H.*

Proof. Since the VC-dimension of both balls and half-spaces in $\mathbb{R}^d$ is $d + 1$, the VC-dimension of the set of intersections of balls and up to two half-spaces is $O(d \ln d)$. Therefore, by a standard VC-bound [Vapnik and Chervonenkis, 1971], if we see an iid sample $\mathfrak{X}$ of size $n = O(\frac{1}{\gamma^2}(\ln^2 \frac{d}{\gamma} + \ln \frac{1}{\delta}))$, then with probability at least $1 - \delta$ the empirical measure of any ball intersected with up to two half-spaces will be within γ of its true probability mass. In other words, the fraction of the sample set $\mathfrak{X}$ that lands in any ball intersected with up to two half-spaces will be within γ of the probability that a sample X drawn from P will land in the same set. For the remainder of the proof, assume that this high-probabilty event holds.

First, we show that every half-ball in the set H is an α-approximation to at least one true hyperplane. Suppose otherwise, then there is a half-ball $B^{1/2} = B^{1/2}(\hat{x}, r, \hat{w})$ with $(\hat{x}, \hat{w}) \in H$ that is not an α approximation to any true hyperplane h_i. The center $\hat{x}$ of the half-ball must belong to the positive region K, since it is one of the sample points. If the half-ball $B^{1/2}$ is contained entirely in the set K, then the probability that a new sample X drawn from P will land in the half-ball $B^{1/2}$ is $p^{1/2}(r)$ and therefore the fraction of samples that landed in the half-ball is at least $p^{1/2}(R/2) - \gamma$. But since $p^{1/2}(r) - \gamma \geq \frac{4}{5}\alpha p^{1/2}(r) > \tau$, this contradicts the half-ball being included in the set H. Otherwise, the half-ball contains at least one point y that does not belong to the set K (i.e., it does not belong to any class). Since $\hat{x}$ is in the set K, there is at least one true hyperplane h_i that separates $\hat{x}$ from y. Since $r = R/2 < R$, every other point y' in the half-ball that does not belong to any class must have the same code word as y (since, by assumption, points outside of K that do not belong to any class must have the same code word when they are closer than R), and therefore must be on the same side of h_i as y. It follows that all points in the half-ball on the same side of h_i as $\hat{x}$ (i.e., those points for which the sign of h_i matches the sign of $h_i(\hat{x})$) belong to the set K. But, since the half-ball is not an α-approximation to h_i, this implies that at least an α fraction of the half-ball's volume must belong to the set K. Therefore, the probability that a new sample x drawn from the data distribution p belongs to the half-ball

can be lower bounded as follows:

$$\mathop{\mathbb{P}}_{x\sim p}(x \in B^{1/2}) \geq c_{\text{lb}} \operatorname{Vol}(B^{1/2} \cap K) = c_{\text{lb}} \operatorname{Vol}(B^{1/2}) \frac{\operatorname{Vol}(B^{1/2} \cap K)}{\operatorname{Vol}(B^{1/2})} \geq \alpha p^{1/2}(r).$$

By the uniform convergence argument, the fraction of the samples in $\mathfrak{X}$ contained in the half-ball $B^{1/2}$ is at least $\alpha p^{1/2}(r) - \gamma > \tau$, which contradicts the half-ball being in H. In either case we arrived at a contradiction and it follows that every half-ball in H is an α-approximation to at least one true hyperplane h_i.

Finally, we show that the set H will contain at least one half-ball that is an α-approximation to each true hyperplane h_i. Fix any true hyperplane h_i. By assumption, there is a class ℓ and a point x_0 on the decision surface of h_i so that one half-ball of $B(x_0, R)$ with face h_i is is contained in $K_\ell \subset K$ and the other half-ball is disjoint from K. Suppose WLOG that the half-ball on the negative side of h_i is contained in K (the case when the half-ball on the positive side is contained in K is identical). Define $\rho > 0$ to be the width such that the probability that a new sample X from P lands in the slice of the ball $S = \{x \in B(\hat{x}_0, r) \,:\, h_i(x) \in [-\rho, 0]\}$ is equal to $\tau - \gamma$. Note that, since the half-ball on the negative side of h_i is a subset of K and $\tau - \gamma = \frac{3}{10}\alpha p^{1/2}(r) < p^{1/2}(r)$, such a value of ρ always exists. Since $\tau - \gamma > \gamma$, the uniform convergence argument guarantees that there will be at least one sample point in the slice, say $\hat{x} \in \mathfrak{X}$. Since $\hat{x}$ is within distance $r = R/2$ of the point x_0, the ball $B(\hat{x}, r)$ is contained in $B(x_0, R)$. Therefore, the ball of radius r centered at $\hat{x}$ only contains points that either belong to class ℓ or no class, since only the linear separator h_i passes through this ball. By construction, the half-ball $B^{1/2}(\hat{x}, r, w_i)$ (where w_i is the coefficient vector defining $h_i(x) = w_i^\top x - b_i$) with face parallel to h_i intersects the set K in a slice of width at most ρ and therefore has probability mass at at most $\tau - \gamma$. It follows that the direction $\hat{w}$ that minimizes the number of samples in the half-ball $B^{1/2}(\hat{x}, r, \hat{w})$ will result in the half-ball containing at most a τ fraction of the sample set, and therefore the pair $(\hat{x}, \hat{w})$ will be included in H, and this will be an α-approximation to h_i. □

Naturally, if a half-ball $B^{1/2}(\hat{x}, r, \hat{w})$ is an α-approximation to the linear function h, we expect that the decision surface of $\hat{h}(x) = \hat{w}^\top(x - \hat{x})$ is similar to the decision surface of h. In turn, this suggests that either $\hat{h}(x)$ or $-\hat{h}(x)$ should take similar function values to $h(x)$ (since the coefficient vectors are normalized). We first give a simple probability lemma that bounds the fraction of a ball contained between two parallel hyperplanes, one passing through the ball's center.

Lemma 6. *Let $r > 0$ be any radius and X be a random sample drawn uniformly from the ball of radius r centered at the origin. For any width $0 \leq \rho \leq r/\sqrt{2}$, the probability that the first coordinate of X lands in $[0, \rho]$ can be bounded as follows:*

$$\sqrt{\frac{d}{2^d \pi}} \frac{\rho}{r} \leq \mathop{\mathbb{P}}_{X \sim B(r,0)}(X_1 \in [0, \rho]) \leq \sqrt{\frac{d+1}{2\pi}} \frac{\rho}{r}.$$

Proof. Let B be the ball of radius r centered at the origin and $S = \{x \in B \,:\, x_1 \in [0, \rho]\}$ be the slice of B for which the first coordinate is in the interval $[0, \rho]$. The probability that a uniformly random sample from B lands in the subset S is given by $\operatorname{Vol}(S)/\operatorname{Vol}(B)$, where Vol denotes the (Lebesgue) volume of a set.

We bound the volume of the set S by writing the volume as a double integral over the first coordinate x_1 and the remaining $d-1$ coordinates x_R.

$$\operatorname{Vol}(S) = \int_0^\rho \int_{\mathbb{R}^{d-1}} \mathbb{I}\{\|x_R\|_2^2 \leq r^2 - x_1^2\} \, dx_R \, dx_1$$

Noticing that the inner integral is actually the volume of a $d-1$ dimensional ball of radius $\sqrt{r^2 - x_1^2}$, and using the fact that for any d, the volume of a d-dimensional ball of radius r is $r^d v_d$, where v_d is the volume of the d-dimensional unit ball, we have

$$\operatorname{Vol}(S) = v_{d-1} \int_0^\rho (r^2 - x_1^2)^{(d-1)/2} \, dx_1.$$

Upper bounding the integrand by r^{d-1} gives that $\operatorname{Vol}(S) \leq v_{d-1} \rho r^{d-1}$. Lower bounding the integrand by $(r^2 - \rho^2)^{(d-1)/2}$ and using the fact that $\rho \leq \frac{r}{\sqrt{2}}$ we have that $\operatorname{Vol}(S) \geq v_{d-1} \frac{1}{\sqrt{2^{d-1}}} \rho r^{d-1}$. Dividing both inequalities by the

volume of B, which is $r^d v_d$, and using the fact that for all d we have $\frac{v_{d-1}}{v_d} \in [\sqrt{\frac{d}{2\pi}}, \sqrt{\frac{d+1}{2\pi}}]$ gives

$$\sqrt{\frac{d}{2^d\pi}}\frac{\rho}{r} \leq \mathop{\mathbb{P}}_{X\sim B}(X \in S) \leq \sqrt{\frac{d+1}{2\pi}}\frac{\rho}{r},$$

as required. □

Using Lemma 6, we show the following:

Lemma 7. *Let the half-ball $B^{1/2}(\hat{x}, r, \hat{w})$ be an α-approximation to the linear function $h(x) = w^\top x - b$ with $\|w\| = 1$, $\hat{x} \in \mathcal{X}$, and $\alpha < \frac{1}{2}$. Let D be the diameter of $\mathcal{X}$. If $h(\hat{x}) < 0$ then for all $x \in \mathcal{X}$ we have*

$$|h(x) - \hat{h}(x)| \leq \left(2D + \sqrt{\frac{2^d\pi}{d}}\frac{r}{2}\right)\sqrt{\alpha},$$

where $\hat{h}(x) = \hat{w}^\top(x - \hat{x})$. Otherwise, if $h(\hat{x}) > 0$ then the same upper bound holds for $|h(x) + \hat{h}(x)|$.

Proof. Suppose that $h(\hat{x}) < 0$ and let X be a uniformly random sample from the half-ball $B^{1/2} = B^{1/2}(\hat{x}, r, \hat{w})$. By assumption, we know that $\mathbb{P}(h(X) < 0) \leq \alpha$.

First we show that $\|w - \hat{w}\|$ is small. Since $\alpha < 1/2$ we have that $w^\top\hat{w} > 0$. To see this, notice that we must have $h(\hat{x} + r\hat{w}) \geq 0$, since otherwise at least half of the half-ball would be on the negative side of h. Define $g(x) = w^\top(x - \hat{x})$ to be the linear function whose decision surface runs parallel to that of h but passes through the point $\hat{x}$. Since $h(x) = g(x) + h(\hat{x}) \leq g(x)$, we have that $\alpha > \mathbb{P}(h(X) < 0) \geq \mathbb{P}(g(X) < 0)$. Moreover, since the decision surface of g passes through the center of the half-ball $B^{1/2}$ and the uniform distribution on the half-ball is radially symmetric about the point $\hat{x}$, we have that $\mathbb{P}(g(X) < 0) = \frac{\theta(w,\hat{w})}{\pi}$. It follows that $\theta(w, \hat{w}) \leq \pi\alpha$. Using this fact, we can bound $\|w - \hat{w}\|$ as follows:

$$\|w - \hat{w}\|^2 = \|w\|^2 + \|\hat{w}\|^2 - 2w^\top\hat{w} = 2(1 - w^\top\hat{w}).$$

Since $w^\top\hat{w} = \cos(\theta(w, \hat{w}))$ and on the interval $[0, \pi/2]$, the $\cos(\theta)$ function is decreasing and lower bounded by $1 - \frac{2}{\pi}\theta$, we have that $2(1 - w^\top\hat{w}) \leq 4\alpha$. Taking the square root gives that $\|w - \hat{w}\| \leq 2\sqrt{\alpha}$.

Next we show that $|h(\hat{x})|$ (the distance from $\hat{x}$ to the decision surface of h) is not too large. The half-ball $B^{1/2}(\hat{x}, r, w)$, whose direction w matches the coefficient vector of h is one half-ball centered at $\hat{x}$ of radius r minimizing the fraction of its volume contained on the same side of h as $\hat{x}$. This is because every point in the ball $B(\hat{x}, r)$ not on the same side as $\hat{x}$ is contained in $B^{1/2}(\hat{x}, r, w)$. Let Y be uniformly sampled from $B^{1/2}(\hat{x}, r, w)$. By construction of the half-ball Y is sampled from, we have that $\mathbb{P}(h(X) < 0) \geq \mathbb{P}(h(Y) < 0)$, which gives

$$\alpha \geq \mathop{\mathbb{P}}_{X\sim B^{1/2}(\hat{x},r,\hat{w})}\big(h(X) < 0\big) \geq \mathop{\mathbb{P}}_{Y\sim B^{1/2}(\hat{x},r,w)}\big(h(Y) < 0\big) \geq \sqrt{\frac{d}{2^d\pi}}\frac{2|h(\hat{x})|}{r},$$

which implies that

$$|h(\hat{x})| \leq \sqrt{\frac{2^d\pi}{d}}\frac{r\alpha}{2}.$$

Finally, let x' be any point on the decision surface of h, so that $h(x) = w^\top(x - x')$. Combining the above calculations we have

$$\begin{aligned}
|h(x) - \hat{h}(x)| &= |w^\top(x - x') - \hat{w}^\top(x - \hat{x})| \\
&= |w^\top(x - \hat{x}) + w^\top(\hat{x} - x') - \hat{w}^\top(x - \hat{x})| \\
&= |(w - \hat{w})(x - \hat{x}) + w^\top(\hat{x} - x')| \\
&\leq \|w - \hat{w}\|\|x - \hat{x}\| + |h(\hat{x})| \\
&\leq 2\sqrt{\alpha}D + \sqrt{\frac{2^d\pi}{d}}\frac{r\alpha}{2} \\
&\leq (2D + \sqrt{\frac{2^d\pi}{d}}\frac{r}{2})\sqrt{\alpha},
\end{aligned}$$

as required. The proof of the case when $h(x) > 0$ follows by applying the above arguments to the function $-h$. $\square$

Recall that for any hyperplane $h(x) = w^\top x - b$ with $\|w\|_2 = 1$, the distance from point x to the decision surface of h is $|h(x)|$. The above lemma implies that if $B^{1/2}(\hat{x}, r, \hat{w})$ is an α-approximation to h, then either $\hat{h}$ or $-\hat{h}$ will have the same sign as h for all points in $\mathcal{X}$ except those in a margin of width $O(\sqrt{\alpha})$ around h. Under the uniform distribution on K, the probability mass of the margins surrounding the true hyperplanes isn't large, which results in low error for the classification rule.

Theorem 4. *Suppose there exists a consistent output code classifier satisfying the boundary features condition with scale parameter R. For any desired error rate $\epsilon > 0$ and confidence parameter $\delta > 0$, let $\hat{f}$ be the output classifier of Algorithm 4 run with parameters $r = R/2$ and $\tau = \frac{1}{2}\alpha p^{1/2}(R/2)$ on a sample of size $n = O(\frac{1}{\gamma^2}(\ln^2 \frac{d}{\gamma} + \ln \frac{1}{\delta}))$ where $\gamma = \frac{1}{5}\alpha p^{1/2}(R/2)$ and $\alpha = \Omega\big(\epsilon^2/(m^2 2^d R^2 D^{2d} v_d^2 c_{\mathrm{lb}}^2)\big)$, and D is the diameter of $\mathcal{X}$ (giving $\gamma = \Omega\big((\frac{\epsilon}{m v_d c_{\mathrm{ub}}})^2 (\frac{R}{4D})^d\big)$). Then with probability at least $1 - \delta$, the error of $\hat{f}$ is at most ϵ.*

Proof. By Lemma 5, for the parameter settings τ and r and the given sample size, with probability at least $1 - \delta$ every half-ball included in the set H will be an α-approximation to some true hyperplane h_i, and every true hyperplane h_i is α-approximated by at least one half-ball in H. Assume that this high probability event occurs.

Let $H = \{(\hat{x}_1, \hat{w}_1), \ldots, (\hat{x}_M, \hat{w}_M)\}$ be the set of of half-balls produced by the algorithm and define $\hat{h}_i(x) = \hat{w}_i^\top (x - \hat{x}_i)$ for $i = 1, \ldots, M$ to be the corresponding linear functions. Algorithm 4 uses these hyperplanes to partition the space $\mathcal{X}$ into a collection of polygonal regions and assigns a unique class label to each cell in the partition. Notice that negating any of the $\hat{h}_i$ functions does not change the partitioning of the space. Therefore, negating any subset of the $\hat{h}_i$ will not change the permutation-invariant error of the resulting classifier.

Let $I_1, \ldots, I_m$ be a partition of the set of indices $\{1, \ldots, M\}$ such that for all $j \in I_i$, we have that $B^{1/2}(\hat{x}_j, \hat{w}_j, r)$ is an α-approximation to h_i. By Lemma 7, we know that for at least one $g \in \{\hat{h}_j, -\hat{h}_j\}$, we have that

$$|h_i(x) - g(x)| \leq \left(2D + \sqrt{\frac{2^d \pi}{d}} \frac{r}{2}\right) \sqrt{\alpha}$$

Since negating any of the functions $\hat{h}_j$ does not change the error of the resulting classifier, assume WLOG that the above holds for $g = \hat{h}_j$.

This implies that whenever $|h_i(x)| > c\sqrt{\alpha}$, where $c = 2D + \sqrt{\frac{2^d \pi}{d}} \frac{R}{4}$, then for every $j \in I_i$, the sign of $\hat{h}_j(x)$ is the same as the sign of $h_i(x)$. It follows that for points that are not within a margin of $c\sqrt{\alpha}$ of any of the true hyperplanes, every $\hat{h}_j$ function with $j \in I_i$ will have the same sign as h_i for all $i = 1, \ldots, m$. It follows that the classifier can only error on points that are within a $c\sqrt{\alpha}$ margin of one of the true hyperplanes.

Using Lemma 6 we can bound the probability that a sample X drawn uniformly from K lands in the $c\sqrt{\alpha}$-margin of hyperplane h_i as follows:

$$\mathbb{P}(X \text{ in } c\sqrt{\alpha}\text{-margin of } h_i) \leq 2\sqrt{\frac{d+1}{2\pi}} \frac{c\sqrt{\alpha}}{D} D^d v_d c_{\mathrm{ub}},$$

where D is the diameter of X. We can make this upper bound equal to ϵ/m by setting

$$\alpha = \frac{\pi}{2(d+1)} \left(\frac{\epsilon D}{m c D^d v_d c_{\mathrm{lb}}}\right)^2 = \Omega\left(\frac{\epsilon^2}{m^2 2^d R^2 D^{2d} v_d^2 c_{\mathrm{lb}}^2}\right)$$

Applying the union bound to the m hyperplanes $h_1, \ldots, h_m$ shows that the error of $\hat{f}$ is at most ϵ.

$\square$

It is worth noting that if we set τ to be smaller than necessary in Theorem 4 and set $\gamma = \frac{2}{5}\tau$, then the result still holds. In other words, setting the parameter τ to be conservatively small based on estimates of $\mathrm{Vol}(K)$ and R is sufficient to achieve high accuracy.

7 Conclusion

In this work we showed how to exploit the implicit geometric assumptions made by output code techniques under the well studied cases of one-vs-all and well separated codewords, and for a novel boundary features condition that captures the intuition that every semantic feature should be significant. We provide both unsupervised and active learning algorithms capable of handling learning when the data distribution is nearly uniform and under the much weaker condition that the data density has thick level sets. In all cases, our algorithms show that the implicit assumptions of output code learning can be used to learn from very limited labeled data.

Acknowledgments

This work was supported in part by NSF grants CCF-1422910, CCF-1535967, a Sloan Research Fellowship, a Microsoft Research Faculty Fellowship, and a Google Research Award.

References

K. Balasubramanian, P. Donmez, and G. Lebanon. Unsupervised supervised learning ii: Margin-based classification without labels. In *AISTATS*, pages 137–145, 2011a.

K. Balasubramanian, P. Donmez, and G. Lebanon. Unsupervised supervised learning ii: Margin-based classification without labels. In *Journal of Machine Learning Research*, volume 12, pages 3119–3145, 2011b.

M.-F. Balcan, A. Blum, and Y. Mansour. Exploiting ontology structures and unlabeled data for learning. In *Proceedings of the 31st International Conference on Machine Learning (ICML)*, pages 1112–1120, 2013.

A. Beygelzimer, J. Langford, and P. Ravikumar. Solving multiclass learning problems via error-correcting output codes. *ALT 2009*, 2009.

A. Blum and S. Chawla. Learning from labeled and unlabeled data using graph mincuts. In *Proceedings of the 18th International Conference on Machine Learning (ICML)*, pages 19–26, 2001.

K. Chaudhuri and S Dasgupta. Rates of convergence for the cluster tree. In *Advances in Neural Information Processing 23 (NIPS)*, pages 343–351, 2010.

T. G. Dietterich and G. Bakiri. Solving multiclass learning problems via error-correcting output codes. *Journal of Artificial Intelligence Research*, pages 263–286, 1995.

P. Donmez, G. Lebanon, and K. Balasubramanian. Unsupervised supervised learning i: Estimating classification and regression errors without labels. In *Journal of Machine Learning Research*, volume 11, pages 1323–1351, 2010.

J. Langford and A. Beygelzimer. Sensitive error correcting output codes. *COLT*, 2005.

T. Mitchell, W. Cohen, E. Hruschka, P. Talukdar, J. Betteridge, A. Carlson, B. Dalvi, M. Gardner, B. Kisiel, J. Krishnamurthy, N. Lao, K. Mazaitis, T. Mohamed, N. Nakashole, E. Platanios, A. Ritter, M. Samadi, B. Settles, R. Wang, D. Wijaya, A. Gupta, X. Chen, A. Saparov, M. Greaves, and J. Welling. Never-ending learning. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI-15)*, 2015.

M. Mohri, A. Rostamizadeh, and A. Talwalkar. *Foundations of Machine Learning*. MIT press, 2012.

M. Palatucci, D. Pomerleau, G. Hinton, and T. Mitchell. Zero-shot learning with semantic output codes. In *Neural Information Processing Systems*, 2009.

J. Steinhardt and P. Liang. Unsupervised risk estimation with only structural assumptions. 2016. (Preprint from `http://cs.stanford.edu/~jsteinhardt/publications/risk-estimation/preprint.pdf`).

I. Steinwart. Fully adaptive density based clustering. In *Annals of Statistics*, volume 43, pages 2132–2167, 2015.

S. Thrun. *Explanation-Based Neural Network Learning: A Lifelong Learning Approach*. Kluwer Academic Publishers, Boston, MA, 1996.

S. Thrun and T. Mitchell. Learning one more thing. In *Proc. 14th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1217–1225, 1995a.

Sebastian Thrun and Tom M. Mitchell. Lifelong robot learning. *Robotics and Autonomous Systems*, 15(1-2):25–46, 1995b.

V. Vapnik and A. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16(2):264–280, 1971.

8 Appendix for Error Correcting Output Codes

The following two Lemma's show that when there exists a consistent linear output code whose code words are well separated and the linear separators are in general position, then there must be a gap $g > 0$ between the classes.

Lemma 8. *Let $i \neq j$ be any pair of classes whose codewords disagree on the k^{th} bit. Then for any points $x \in K_i$ and $y \in K_j$, the line segment $[x, y]$ intersects with the line $h_k = 0$.*

Proof. Without loss of generality, suppose that $C_{ik} = 1$ and $C_{jk} = -1$. Then, from the definition of K_i and K_j, we have that $h_k(x) > 0$ and $h_k(y) < 0$. The function $f(t) = h_k((1-t)x + ty)$ is continuous and satisfies $f(0) = h_k(x) > 0$ and $f(1) = h_k(y) < 0$. It follows that there must be some $t_0 \in (0, 1)$ such that $f(t_0) = 0$. But this implies that the point $z = (1 - t_0)x + t_0 y \in [x, y]$ satisfies $h_k(z) = 0$ and it follows that $h_k = 0$ intersects with $[x, y]$ at the point z. □

Next, we show that when the consistent linear output code makes at most β errors when predicting the code word of a new example and the Hamming distance of the code words is at least $2\beta + d + 1$, then there must be a minimum gap $g > 0$ between any pair of points belonging to different classes.

Lemma 9. *Suppose there exists a consistent linear output code classifier that never makes more than β mistakes when predicting the code word for a new example and whose code words have Hamming distance at least $2\beta + d + 1$. There exists a distance $g > 0$ such that $\|x - y\| \geq g$ whenever points x and y belong to different classes.*

Proof. For sets A and B, let $d(A, B) = \min_{a \in A, b \in B} \|a - b\|$ denote the distance between them and recall that for each $i = 1, \ldots, L$, we defined $K_i = \{x \in \mathcal{X} : d_{\text{Ham}}(h(x), C_i) \leq \beta\}$ to be the set of points that belong to class i.

Fix any pair of classes i and j and suppose for contradiction that $d(K_i, K_j) = 0$. This implies that there are two code words $c, c' \in \{\pm 1\}^m$ such that $d_{\text{Ham}}(c, C_i) \leq \beta$, $d_{\text{Ham}}(c', C_j) \leq \beta$, and the distance between $A = \{x \in \mathcal{X} : h(x) = c\}$ and $B = \{x \in \mathcal{X} : h(x) = c'\}$ is 0. First, we construct a point x that belongs to $\overline{A} \cap \overline{B}$, where $\overline{A}$ and $\overline{B}$ denote the closure of A and B, respectively. Since $d(A, B) = 0$, there exists a sequence of points $x_1, x_2, \ldots \in A$ such that $d(x_n, B) \to 0$ as $n \to \infty$. But, since A is bounded, so is the sequence (x_n), and therefore by the Bolzano-Weierstrass theorem, (x_n) has a convergent subsequence. Without loss of generality, suppose that (x_n) itself converges to the point x. Then x is a limit point of A and therefore belongs to the closure of A. On the other hand, since the function $z \mapsto d(z, B)$ is continuous, it follows that $d(x, B) = \lim_{n \to \infty} d(x_n, B) = 0$ and therefore x is also in the closure of B.

Now let k be any index such that the code words c and c' differ on the k^{th} entry. Next, we show that $h_k(x) = 0$. For each integer $n > 0$, let $C_n = B(x, 1/n)$ be the ball of radius $1/n$ centered at x. Since x belongs to the closure of A and C_n is a neighborhood of x, we can find some point, say x_n that belongs to the intersection $A \cap C_n$. Similarly, we can find a point y_n belonging to $B \cap C_n$. Since the line segment $[x_n, y_n]$ passes from A to B, Lemma 8 guarantees that there is a point $z_n \in [x_n, y_n] \subset C_n$ such that $h_k(z_n) = 0$. But, by construction, the sequence z_n is converging to x and, since linear functions are continuous, it follows that $h_k(x) = \lim_{n \to \infty} h_k(z_n) = 0$.

But this leads to a contradiction: since the codewords c and c' must disagree on at least $d + 1$ entries, at least $d + 1$ of the linear separators $h_1, \ldots, h_m$ intersect at the point x, which contradicts our assumption that at most d lines intersect at any point $x \in \mathcal{X}$. Therefore, we must have $d(K_i, K_j) > 0$. Since there are finitely many classes, taking $g = \min_{i,j} d(K_i, K_j)$ completes the proof. □

8.1 Using the Thick Level Set Condition Instead of Round Sets

The following result is an anaysis of Algorithm 1 analogous to Theorem 2 that uses the thick level set condition instead of requiring that the positive regions are round.

Theorem 5. *Suppose there exists a consistent output code classifier that makes at most β mistakes when predicting the codeword of a new example and that p is the uniform distribution over the positive region $K = \bigcup_{i=1}^{L} K_i$ with thick level sets up to level $\lambda_0 = \frac{1}{\text{Vol}\, K}$ with constants C and σ_0. Let $\sigma = \min\left(\frac{g}{(2C+6)}, \sigma_0\right)$ and run Algorithm 1 on a sample of size*

$$n \geq \frac{\text{Vol}(K)}{\epsilon \sigma^d v_d} \left(\ln \frac{2^d \,\text{Vol}(K)}{\epsilon \sigma^d v_d} + \ln \frac{1}{\delta} \right).$$

Then with probability at least $1 - \delta$, the output classifier has zero error.

Proof. For any $\lambda \in [0, \frac{1}{\mathrm{Vol}(K)}]$, the λ level set of p is exactly equal to K and has L connected components, given by the positive regions $K_1, \ldots, K_L$. By Lemma 4, with probability at least $1 - \delta$, the single linkage graph for our sample at connection radius $r_c = (2C + 6)\sigma$ will output one large cluster for each of the connected components $K_1, \ldots, K_L$. With probability 1, a new sample x will land in the set K, which means that it will belong to one of the connected components K_i. In such a case, the nearest cluster to x will be the cluster in K_i, and therefore x is classified correctly. $\square$